


CUSTOMER NO. 23932

Docket No.: 50886-00024USPX
(PATENT)



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of:
Mathieu Robart

Application No.: 10/811481

Confirmation No.: 8876

Filed: March 26, 2004

Art Unit: 2671

For: COMPUTER GRAPHICS

Examiner: Not Yet Assigned

CLAIM FOR PRIORITY AND SUBMISSION OF DOCUMENTS

MS MISSING PARTS
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

I hereby certify that this correspondence is being deposited with the U.S. Postal Service with sufficient postage as First Class Mail, in an envelope addressed to: MS MISSING PARTS, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date shown below.

Dated: 08-4-04

Signature: M. Szuwalski
(Margo Barbarash)

Dear Sir:

Applicant hereby claims priority under 35 U.S.C. 119 based on the following prior foreign application filed in the following foreign country on the date indicated:

Country	Application No.	Date
European Community	03252037.4	March 31, 2003

In support of this claim, a certified copy of the said original foreign application is filed herewith.

Dated: 8/3/04

Respectfully submitted,

By

Andre M. Szuwalski

Registration No.: 35,701

JENKENS & GILCHRIST, A PROFESSIONAL CORPORATION

1445 Ross Avenue, Suite 3200

Dallas, Texas 75202

(214) 855-4500

Attorneys For Applicant

THIS PAGE BLANK (USPTO)



**Europäisches
Patentamt**

**European
Patent Office**

**Office européen
des brevets**

Bescheinigung

Certificate

Attestation

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

03252037.1

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk

THIS PAGE BLANK (USPTO)



Anmeldung Nr:
Application no.: 03252037.1
Demande no:

Anmelde tag:
Date of filing: 31.03.03
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

STMicroelectronics, Ltd.
1000 Aztec West
Almondsbury,
Bristol, BS32 4SQ
GRANDE BRETAGNE

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se referer à la description.)

Computer graphics

In Anspruch genommene Priorität(en) / Priority(ies) claimed /Priorité(s)
revendiquée(s)
Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

G06F15/60

Am Anmelde tag benannte Vertragstaaten/Contracting states designated at date of
filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IT LU MC NL
PT SE SI SK TR LI

THIS PAGE BLANK (USPTO)

COMPUTER GRAPHICS

The present invention is concerned with computer graphics and in particular, but not exclusively, with rendering an image including shadows in a graphics scene generated by a computer.

The field of computer graphics is a fast changing one, in which trade-offs are made between the quality of the images being produced and the processing power of the computing platform being used. The trade-off becomes even more acute when the images need to be displayed in real-time, for example in the video-gaming industry.

The rendering of an image on a computer screen entails vast amounts of processing. For example, it is necessary to have a so-called "global scene" (also known as the "world scene") of the images to be displayed, and which will be hereinafter referred to as the scene. Broadly speaking, a scene can be thought of as a snap-shot or picture of the image to be displayed on the screen at any instant in time. As would be expected, a scene will itself comprise many different objects each having their own geometry. For example, a global scene might be the interior of a particular room in a house. The room might have windows, furniture, a TV, etc. Each object, for example TV, table, window will have a different geometry and will need to be created with the correct dimensions and co-ordinates on the screen in relation to the other images.

These objects are defined in three dimensions (3D), but will be rendered onto a two-dimensional (2D) computer screen. The technique for rendering a 3D object onto a 2D display involves firstly breaking down the 3D object into polygons defined by primitives. A popular primitive used is a triangle having three vertices. Thus, a 3D image can be transformed into a plurality of triangles each being defined by a unique set of vertices where each vertex would typically contain information relating to co-ordinates (x, y, z), colour, texture and lighting. It should

be understood that a fairly large storage area is needed to accommodate the vertex information.

Broadly speaking, the creation of an image on a computer graphics screen can be thought of as consisting of a geometry stage and a rendering stage. In existing systems, the geometry stage is responsible for transformation and lighting, in which the 3D object is converted into a number of polygons defined by a set of suitable primitives. Consider an interactive computer game where the user controls the motion of his player, as the player moves forward or backward the objects in the frame will need to be transformed so that they appear closer to and further away from the user, respectively.

In the rendering stage the transformed vertices are placed in a frame buffer in digital form. The frame buffer can in fact comprise a number of buffers, e.g. a colour buffer, depth buffer, stencil buffer, accumulation buffer. The frame buffer needs to be continuously managed and updated as the frame (or scene) changes. In the rendering stage the pixels are drawn and the frame buffer stores lighting, colour and intensity information for each pixel that will be enabled. The digital frame data is then converted to an analogue signal to be used to provide the final image for the actual 2D computer display.

The lighting of a scene, and of the objects within it, contributes largely to the overall success of the scene when displayed to the user. There are various methods that have been proposed for dealing with lighting.

The so-called "Shadow Map" method is a popular shadowing technique for determining whether surfaces in a scene are in shadow or illuminated. See Figure 1. The method is based on the idea of a two-pass rendering of the scene. In a first pass, the scene is rendered from the position of the light source 2. A depth buffer (the so-called classic zbuffer) stores the depth representing the distance to the object 4 closest to the light source 2 for each pixel. The content of the depth buffer is the actual shadow map 6, which is stored.

Then during a second pass the scene is rendered classically from the viewpoint of the camera 8. For this pass, the depth buffer holds as depths distances from the camera. To ensure that polygons are drawn correctly on the screen, each pixel of each polygon as it is considered is depth compared with any pixel already at a corresponding pixel location in the frame buffer. In one such arrangement, the depth buffer is the same size as the frame buffer, and is used to maintain depth data in the form of a depth value for each pixel that has been written to the frame buffer. When a new pixel is being considered, its depth value is compared with the depth value associated with any pixel that has already been written to the frame buffer at the new pixel's location. In the event that the new pixel is behind the old, then it is discarded because it will be obscured. Conversely, if the new pixel is in front of the old, the new pixel's depth value replaces that of the old pixel in the depth-buffer, and colour data is retrieved from associated texture memory and written over the old pixel's colour data in the frame buffer.

For each rasterised point, its world-space coordinates are generated (if not already stored) and converted into the light coordinate space. Its distance to the light (or depth) is compared with the corresponding depth stored in the shadow map. The result of the comparison will determine if the point is in shadow (P0) or illuminated (P1).

Let us consider the two points P0, P1 viewed from the camera. Using an inverse ModelView transformation, the world coordinates of P0 and P1 can be determined, and then transformed again in the light coordinate system. For each point, its distance to the light is compared to the depth stored in the corresponding pixel of the shadow map:

- If the distance to the light is greater than the corresponding depth, the point is in shadow.
- If the distance to the light is "near" the corresponding depth, the point is lit.

- The distance to the light cannot be lower than the corresponding depth in theory.

In this example, d_1 is the distance between P_1 and the light. By construction, the depth of P_1 is also equal to d_1 and consequently P_1 is in the light. In the other case, d_0 (the distance between P_0 and the light) is greater than d_0' (the depth of P_0). Consequently, P_0 is in shadow.

However, the disadvantage of the shadow map method is that it only provides information as to whether a point is in light or shadow and does not cater for soft shadows (i.e. differing shadow intensities).

It is an aim of embodiments of the present invention to allow soft shadows to be rendered, without hugely increasing the processing overhead.

According to one aspect of the invention there is provided a method of rendering an image comprising at least one light source, a first, shadow-casting object with a second, shadow-receiving object located on the side of the first shadow-casting object remote from said at least one light source, the method comprising: generating a shadow mask which identifies for each of a plurality of pixels on the shadow receiving surface a grey level representing the intensity of shadow in each pixel, the intensity being determined utilising the distance between the shadow-casting object and the shadow-receiving object.

According to another aspect of the invention there is provided a computer system for rendering an image comprising at least one light source, a first shadow-casting object with a second shadow-receiving object located on the side of the first shadow-casting object remote from said at least one light source, the computer system comprising: a shadow mask memory in which is stored a shadow mask which identifies for each of a plurality of pixels on the shadow receiving object a grey level representing the intensity of shadow in each pixel, the intensity having been determined utilising the distance between the shadow-casting object and

the shadow-receiving object; and processing means for utilising the shadow mask to render the image by adjusting the texture of each pixel based on said grey level representing the intensity of shadow.

Another aspect of the invention provides an image rendering pipeline including a polygon identification stage and a pixel rendering stage, wherein the pixel rendering stage comprises: a rasteriser which rasterises pixel parameters for each pixel, including a colour parameter; a texture mapping stage which modifies said colour parameter according to texture values; a shadow mask determination stage which generates a shadow mask identifying for each of a plurality of pixels on a shadow receiving surface a grey level representing the intensity of shadow in each pixel, the intensity having been determined utilising the distance between a shadow-casting object and the shadow-receiving object; wherein the texture value is modulated using each grey level whereby soft shadows in the final image can be rendered.

Thus, the present invention makes use of the fact that the intensity of a cast shadow depends on the distance between the object that casts the shadow and the surface in shadow. In existing systems, no information about this geometry is retained at pixel level, but according to the following described embodiments of the present invention this distance information is deduced and is used by the rasteriser for rendering the image.

For a better understanding of the present invention and to show how the same may be carried into effect, reference will now be made by way of example to the accompanying drawings, in which:

Figure 1 is a graphics image illustrating shadow-map geometry;

Figure 2 is a simplified two-dimensional diagram denoting primary and secondary depths;

Figure 3 is a geometric diagram illustrating intersection distances and shadows;

Figure 4 is a diagram illustrating a hierarchy of textures;

Figures 5A and 5B are images using a known shadow-map technique and a shadow-mask technique according to one embodiment of the invention;

Figure 6 is a schematic block diagram of a classic graphics pipeline;

Figure 7 is a schematic diagram of the parts of the pipeline which have been modified in accordance with an embodiment of the invention; and

Figure 8 is a flow chart illustrating a method of rendering an image in accordance with one embodiment of the invention.

The following described embodiment of the invention rests on a technique for rendering an image which gives better information about the illumination of each pixel, according to the geometry of its projected area in the scene and the nature of the corresponding surface. Thus, instead of merely being able to tell whether a given point is in light or in shadow, the intensity of the shadow can be deduced and used to render a more realistic image, by allowing "soft shadows" to be rendered. The intensity of the shadow depends on the distance between the "shadow caster" and the "shadow receiver". The shadow caster is considered to be the first object in the light path from the light source under consideration. The shadow receiver is considered to be the next object in that path, that is the one which receives the shadow. This is sometimes referred to herein as the immediate second occlusion. A simplistic two dimensional view of this is shown in Figure 2. In Figure 2, d_p represents a primary depth, that is the distance to the object 4 closest to the light source. d_s represents a secondary depth, that is the distance of the shadow receiver 10 from the light source. z denotes the axis along which the depth values are taken. A primary depth buffer (232 – Figure7) acts as a classic zbuffer. That is, it stores the depth representing the distance to the object closest to the light source (denoted by reference numeral 4 in Figure 2).

A secondary depth buffer (233) stores the depth of the immediate second occlusion (10), following the algorithm:

```

if (z<=_primary_depth_buffer(x,y))
{
    secondary_depth_buffer(x,y) =
    primary_depth_buffer(x,y) ;
    primary_depth_buffer(x,y) = z;
}
else
    if(z<=secondary_depth_buffer(x,y))
        secondary_depth_buffer(x,y) = z;

```

In that algorithm, the value `primary_depth_buffer(x,y)` denotes the value of d_p for the pixel at the x,y location. The value `secondary_depth_buffer(x,y)` denotes the value ds for the pixel at the x,y location.

Thus, for values of z up to d_p (that is, points between the light source and the object 4), the primary depth buffer 232 holds the classic z value, and the secondary depth buffer 233 holds the initialisation value, e.g. `MAX_DEPTH`. In this form, it is evident that there is no need to compute the difference between the value stored in the first and second buffers for points lying between the light source 2 and the closest object 4 because the secondary depth has not been written. There is no shadow receiver. However, for values of z exceeding d_p , that is points lying between the object 4 and the immediate second occlusion 10 (shadow receiver), the value which is held in the primary depth buffer remains at d_p while the value that is held in the secondary depth buffer is z , the classic depth. Consequently, the difference Δ between the values held in the primary and secondary depth buffers provides the distance (if any) between the point in the light the object 4 and the shadow that it casts on the shadow receiver 10. This distance characterises the intensity of the shadow and its determination allows soft shadows to be rendered.

Figure 3 represents a simple geometry configuration 3-D, where two triangles, triangle T0 and triangle T1 overlap and cast shadows on a polygon P. In this simple scene, different obstruction cases occur and create shadows:

- 1) triangle T0 casts shadows on the polygon P;
- 2) triangle T0 casts shadows on the triangle T1;
- 3) triangle T1 casts shadows on the polygon P.

In order to evaluate the scene, a maximum distance d_{\max} is defined, with which other distances can be compared. The distance d_{\max} can be specified by a user for example, in order to limit the range of penumbra generated. Alternatively, this value can be automatically defined as the largest difference between the first and second depth buffer values (corresponding to the maximum distance between a shadow caster and the corresponding shadow receiver). A first distance d_0 is computed between an intersection point p_{01} in the plane of the triangle T0 and an intersection point p_{11} in the plane of the polygon P. This distance is relatively close to d_{\max} , and the shadow produced will, as a consequence, be very light. The distance computed is the difference between the depth of p_{11} stored in the secondary depth buffer and the depth of p_{01} stored in the primary depth buffer. The depth of the polygon is treated as the depth of triangles: they are all objects.

A second distance d_1 is computed between a second point of intersection p_{02} on triangle T0 and a first point of intersection p_{12} on triangle T1. This distance d_1 is relatively small in comparison to d_{\max} , and consequently the shadow cast on triangle T1 will be sharp and dark. To compute the second distance d_1 , the distance from the light source to the intersection point p_{12} is held in the secondary depth buffer, and the distance from the light source to the intersection point p_{02} is held in the primary depth buffer.

A third distance d_2 is computed between a point of intersection p_{03} on triangle T1 on a point of intersection p_{13} on the polygon. This distance d_2 is also relatively small, generating a dark shadow. To compute the third distance d_2 , the distance

from the light source to the intersection point p13 is held in the secondary depth buffer and the distance from the light source to the intersection point p03 is held in the primary depth buffer.

Computation of the distances d_1 and d_2 result in the rendering of corresponding dark shadows, even if the distance between the light source and the first intersection point in each case is different. In particular, the distance between the light source and intersection point p02 is smaller than the distance between the light source and the intersection point p03 (considering the light source to be at approximately the distance d_{max} from the plane of the polygon). However, notwithstanding this the shadows cast are dark in both cases. This is the fact in real life when a strong and distant light source such as the sun lights a scene.

The difference Δ is stored as part of the texture information in the frame buffer, such that the texture can then be processed in order to define the degree of shadow for each pixel. During a rendering pass, a simple look-up operation is utilised to render the penumbra, using the texture value as a mask. Moreover, by representing the distance Δ as part of the texture value, the results can be reused in subsequent passes at no cost. The mask constituted by the differences Δ is referred to herein as the shadow mask and its generation is discussed in more detail below.

Shadow mask generation

During the initialisation process, the depth buffers are filled with the same value, MAX_DEPTH which is the maximum value that they can store (for example, 0xFFFFFFFF for 32 bit representation). So, if there are no second objects (that is no shadow receivers), no value is written into the secondary depth buffer and it remains initialised. Consequently, the Δ value is computed only when there is a difference to compute, that is when there are values in the primary and second buffers, different from the initialisation value. For the same reason, the maximum

and minimum and updated when a Δ value is computed. Thus, each cell of the primary and secondary depth buffers are processed together, producing a delta value Δ according to their difference:

- If the corresponding secondary depth buffer cell `secondary_depth_value(x,y)` is not valid (still initialised to `MAX_DEPTH`), then:

$$\text{delta} = \text{MAX_DEPTH}$$

- Else, the difference is stored:

$$\text{Delta} = \text{secondary_depth_value}(x,y) - \text{primary_depth_value}(x,y)$$
- The minimum and maximum value of Δ is updated, only when a **valid difference has been computed, i.e. different from MAX_DEPTH.**

Once the difference between the primary and secondary depth buffers is computed, the shadow mask can be generated. This mask codes the intensity of the shadow in each pixel according to a level of grey. A simple process for "painting" a mask with the correct values is discussed below.

A hierarchy of textures, each representing a point at a different level of "blur" is pre-generated (see Figure 4).

For describing the shadow mask as a texture, the shadow mask itself takes the form of a floating-point memory surface 234, initialised to (1.0, 1.0, 1.0, 1.0). For each of its cells, the delta Δ previously generated determines the corresponding level of blur. The Δ value permits the index of the texture to be determined from the texture hierarchy, the index being used to simulate the level of blur at this pixel location. In the following equation, `Tex_pixel` is an index in the hierarchy of textures extending between `TEXMIN` and `TEXMAX` (see Figure 4).

$$\text{Tex_pixel} = \frac{(\text{delta} - \text{min}) * \text{TEXMAX} + (\text{max} - \text{delta}) * \text{TEXMIN}}{\text{max} - \text{min}}$$

where *min* and *max* are the minimum and maximum distance between a shadow caster and a shadow receiver, and *TEXMIN* and *TEXMAX* are the index of the first and last texture in the hierarchy.

If delta is different from MAX_DEPTH (indicating that there is indeed a shadow to cast), the pixels of the corresponding texture are copied into the shadow mask at the current location, following a blending process:

- A local distance L1 is computed as the difference between the depth at the pixel location (corresponding to the centre of the texture) and the depth at the current pixel of the texture
- A second distance L2 = L1 – delta is computed
- These two distances, compared to a threshold, indicate if the current pixel of the texture is copied into the mask or not.

This algorithm permits to detect any disruption in the distance between the light receiving object 4 (the shadow caster) and the potential object in shadow (the surface 10 of the immediate second occlusion). If such a disruption is detected, then the corresponding points are in a penumbra zone, and its intensity depends on the distance between the shadow caster 4 and the shadow receiver 10.

The approach generates more realistic results than the classic shadow map method. Compare Figure 5A, representing the existing shadow cast method with Figure 5B where a shadow mask as described above is utilised.

Figure 6 illustrates the architecture of a classic graphics pipeline. The architecture is illustrated using a number of connected blocks, each block representing a function in the pipeline. The top row of blocks relates to triangle/vertex processing, and the lower row of blocks relate to pixel/fragment processing. The stages of this pipeline are well understood in the art, and therefore only a brief description follows. In the pipeline, the geometry of the image is defined in a model view transform stage 100, lighting analysis stage 102 and projection analysis stage 104. Primitives, such as polygons (normally made

up of triangles) are assembled in a primitive assembly stage 106. This is followed by a frustum culling stage 108 and a clipping stage 110. This is followed by a perspective divide stage 112 and a back face culling stage 114. Finally, the view port for the scene to be rendered is determined in a view port stage 116. Each primitive is supplied to the pixel/fragment processing pipeline, commencing at a rasteriser stage 118 which rasterises the required pixel parameters, such as depth, colour etc. The output of the rasteriser is supplied to a texture unit 120 which applies texture values to the rasterised pixels. The output of the texture unit is supplied to a fog unit 122 which computes a fog value based on the distance from the camera and applies the fog value to the texture mapped pixels. The texture and fog mapped pixels are then subject to a sequence of tests, including a scissor test 124, an alpha test 126, a stencil test 128 and a depth test 130. The scissor test, alpha test and stencil test are not discussed further herein. The depth test has already been explained earlier in relation to the production of a shadow map for identifying shadows in the image. The depth test is associated with a depth buffer 132 which is the classic zbuffer which has already been referred to. The output of the depth test 130 is applied to a blending unit 134 which carried out a blending operation.

Figure 7 illustrates a modified pipeline to implement the described embodiment of the present invention. In particular, the part circled in Figure 6 is shown in Figure 7 as a modified depth test 230 associated with a primary depth buffer 232 and a secondary depth buffer 233. In addition, the output from the blending unit 134 is applied to an image shader 236. A feedback from the output of the image shader is shown as an input to the texture mapping stage 220 (in the form of the shadow mask), where it is utilised in a fragment shader 238 (described later).

Figure 8 is a flow chart illustrating how the shadow mask can be utilised in a rendering procedure to render an image. Multiple passes are implemented in order to render the final picture, as follows:

a first pass for rendering the shadow map from the current light source;

generation of the shadow mask using the primary and secondary depth buffers;

a second pass for rendering the frame according to this light source from the camera p.o.v. using the shadow mask.

This process is repeated for each light source.

The first step S1 illustrated in Figure 8 is initialisation of the dual depth buffers 232, 233. The primary depth buffer 232 is created in unsigned INT representation, for use as the shadow map buffer. The secondary depth buffer 233 is created as a secondary unsigned INT channel. The buffers are initialised as described above. In the subsequent stage S2, the scene is rendered from the light point of view and the shadow map is generated in the primary depth buffer. At step S3, the shadow mask is generated according to the shadow mask generation technique discussed above. Step S2 constitutes the first rendering pass, and step S3 is carried out after the first rendering pass before the second rendering pass to be described later. After generation of the shadow mask at step S3, the primary depth buffer is initialised for use as the classic zbuffer at step S4. The shadow map which was held in the primary depth buffer is moved to another buffer or memory store 240 so that it is not overwritten, because it is needed in the subsequent, second pass. Next, the camera is defined and the scene is rendered at step S5 from the camera point of view, which constitutes the second rendering pass. At step S5, the following steps are accomplished for each pixel:

The following data is used in the fragment shader of the second pass.

- the shadow map;
- the shadow mask;
- the light model view matrix;
- the light projection matrix;
- the camera model view matrix;

the view port parameters;
 the position of the light;
 the colour generated for the fragment by the rasteriser.

It is looked up in the shadow map to determine whether or not the point is in shadow, if it is in shadow a soft shadow is generated using the shadow mask to modulate the texture value of the pixel, and then a final colour computation is carried out using the colour generated for the fragment by the rasteriser. Finally, at step S6, the result is saved. The shadows, penumbra, colours etc generated by the fragment shader 238 use the shadow mask which has been generated at step S3 to provide soft shadows in the final image, using the textured hierarchy described above.

In one embodiment, in order to render the image, the shadow mask is used in the computation of a form factor in the following equation:

$$\frac{1.0}{\pi(X'^2+Y'^2+1)^2} * \text{ShadowMask}(X, Y)$$

where X and Y are the x,y values of the pixel being rendered. The shadow value parameter of the shadow map (the contents of the primary depth buffer) is read:

`z_shadow=ShadowMap(X,Y)`

and the difference between the depth in the shadow map and the actual difference between the light source and the point P" which represents the position of a point in world space transformed into light space is computed:

$$\text{Delta} = |z - z_shadow|$$

If this difference is lower than a given threshold, the point is in light, otherwise it is in shadow. If the point is in light, the form factor is computed and used to modify the shadow value parameter in the shadow map by shadow_value plus form factor. In this way, the shadow values in the shadow map are modified according to the shadow mask for use in rendering of the image.

THIS PAGE BLANK (USPTO)

THIS PAGE BLANK (USPTO)

CLAIMS:

1. A method of rendering an image comprising at least one light source, a first, shadow-casting object with a second, shadow-receiving object located on the side of the first shadow-casting object remote from said at least one light source, the method comprising:

generating a shadow mask which identifies for each of a plurality of pixels on the shadow receiving surface a grey level representing the intensity of shadow in each pixel, the intensity being determined utilising the distance between the shadow-casting object and the shadow-receiving object.

2. A method according to claim 1, comprising the step of generating the distance between the shadow casting object and the shadow receiving surface by generating a primary depth being the distance between the light source and the shadow-casting object and a secondary depth being the distance between the light source and the shadow-receiving surface, and comparing the primary and secondary depths.

3. A method according to claim 2, wherein the primary depth is stored in a primary depth buffer and the secondary depth is stored in a secondary depth buffer.

4. A method according to any preceding claim, wherein each grey level is generated using a texture hierarchy having respectively different levels of blur.

5. A method according to any preceding claim, wherein a first pass of the image is implemented from the point of view of said at least one light source, prior to generating the shadow mask.

6. A method according to claim 5, wherein a second pass of the image is implemented from the point of view of a camera, said second pass being implemented after the step of generating the shadow mask.
7. A method according to any preceding claim, wherein the shadow mask is used to render the image by adjusting the colour of each pixel in the rendered image based on said grey level representing the intensity of shadow.
8. A computer system for rendering an image comprising at least one light source, a first shadow-casting object with a second shadow-receiving object located on the side of the first shadow-casting object remote from said at least one light source, the computer system comprising:
 - a shadow mask memory in which is stored a shadow mask which identifies for each of a plurality of pixels on the object receiving surface a grey level representing the intensity of shadow in each pixel, the intensity having been determined utilising the distance between the shadow-casting object and the shadow-receiving object; and
 - processing means for utilising the shadow mask to render the image by adjusting the colour of each pixel based on said grey level representing the intensity of shadow.
9. A computer system according to claim 8, which comprises:
 - a primary depth buffer arranged to store a shadow map from a first pass, being for each of the plurality of pixels the distance between said at least one light source and the first shadow-casting object; and
 - a secondary depth buffer arranged to store for each of a plurality of pixels the distance between said at least one light source and the shadow-receiving surface.
10. A computer system according to claim 9, which comprises a shadow map memory into which the shadow map from the primary depth buffer is transferred

for use in a second pass, so that the primary depth buffer can be initialised as a classic depth buffer for that second pass.

11. A computer system according to claim 8, 9 or 10, which comprises shadow mask generation means for generating shadow masks wherein each grey level is generated using a texture hierarchy having respectively different levels of blur.

12. A computer system according to any preceding claim, wherein the processing means includes a fragment shader for utilising the shadow mask to adjust the colour of each pixel in each of a plurality of image fragments.

13. An image rendering pipeline including a polygon identification stage and a pixel rendering stage, wherein the pixel rendering stage comprises:

 a rasteriser which rasterises pixel parameters for each pixel, including a colour parameter;

 a texture mapping stage which modifies said colour parameter according to texture values;

 a shadow mask determination stage which generates a shadow mask identifying for each of a plurality of pixels on a shadow receiving surface a grey level representing the intensity of shadow in each pixel, the intensity having been determined utilising the distance between a shadow-casting object and the shadow-receiving object;

 wherein the texture value is modulated using each grey level whereby soft shadows in the final image can be rendered.

THIS PAGE BLANK (USPTO)

Abstract

COMPUTER GRAPHICS

A method of rendering an image comprising at least one light source, a first, shadow-casting object with a second, shadow-receiving object located on the side of the first shadow-casting object remote from said at least one light source, the method comprising: generating a shadow mask which identifies for each of a plurality of pixels on the shadow receiving surface a grey level representing the intensity of shadow in each pixel, the intensity being determined utilising the distance between the shadow-casting object and the shadow-receiving object.

THIS PAGE BLANK (USPTO)

BEST AVAILABLE COPY




FIG. 1

BEST AVAILABLE COPY




Figure 2




FIG. 3

BEST AVAILABLE COPY

TEXMIN




FIG. 4

TEXMAX

BEST AVAILABLE COPY




Figure 5A

BEST AVAILABLE COPY




FIG. 5B

BEST AVAILABLE COPY




Figure 6





Figure 7



THIS PAGE BLANK (USPTO)